Paper AAS 05-393

LOW-THRUST ORBIT TRANSFER OPTIMIZATION WITH REFINED Q-LAW AND MULTI-OBJECTIVE GENETIC ALGORITHM

Seungwon Lee
Anastassios E. Petropoulos
Paul von Allmen

Jet Propulsion Laboratory
Californian Institute of Technology
4800 Oak Grove Dr. Pasadena, CA

AAS/AIAA Astrodynamics Specialists Conference

Lake Tahoe, CA, August 7-11, 2005

AAS Publications Office, P.O. Box 28130, San Diego, CA 92198

AAS 05-393

LOW-THRUST ORBIT TRANSFER OPTIMIZATION WITH REFINED Q-LAW AND MULTI-OBJECTIVE GENETIC ALGORITHM

Seungwon Lee[1], Anastassios E. Petropoulos[2], and Paul von Allmen[3]

An optimization method for low-thrust orbit transfers around a central body is developed using the Q-law and a multi-objective genetic algorithm. In the hybrid method, the Q-law generates candidate orbit transfers, and the multi-objective genetic algorithm optimizes the Q-law control parameters in order to simultaneously minimize both the consumed propellant mass and flight time of the orbit transfer. Recently, the thrust control condition of the Q-law has been refined by the introduction of the concept of relative effectivity. The refined thrust control condition is tested in the hybrid optimization method, and its contribution to the Q-law performance is analyzed in comparison with the previous thrust control condition given by absolute effectivity. No significant performance difference is found between the Q-law with the absolute effectivity control condition and that with the relative control condition. When the Q-law parameters are optimized, the previously reported shortcomings of the absolute and relative effectivity in the nominal Q-law are mitigated through the optimization process of the other Q-law parameters, which changes the dynamics of the thrust effectivity.

INTRODUCTION

This paper addresses the problem of finding optimal orbit transfers for low-thrust spacecraft. A common goal for the optimization problem is to find the minimum-time, minimum-fuel, or Pareto-optimal trajectory, where the Pareto-optimality means either minimum time for a given fuel or minimum fuel for a given time. In general, these optimization problems are difficult to solve due to the long transfer time and multi-revolutionary transfer.

Various methods have been used to solve this optimization problem. A majority of the work has utilized either direct or indirect techniques.[1] Another quite different approach is to design heuristic control laws.[2,3,4] The advantage of the heuristic control laws is computational efficiency, while the drawback is that the solutions are non-optimal. Recently, the drawback was overcome by combining the heuristic control law

---

[1] Member of the Information and Computer Science Staff at Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA. Email: Seungwon.Lee@jpl.nasa.gov.

[2] Senior Member of the Engineering Staff at Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA. Email: Anastassios.Petropoulos@jpl.nasa.gov.

[3] Senior Member of the Information and Computer Science Staff at Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA. Email: Paul.vonAllmen@jpl.nasa.gov.

with a global optimizer.[5,6] In this hybrid approach, a Lyapunov feedback control law termed Q-law is served as the heuristic control law,[2,3,4] and a multi-objective genetic algorithm is used as a global optimizer for the parameters of the Q-law.[5,6] The synergetic, multi-objective optimization process produces an extended Pareto front with significantly less computational effort than conventional single-objective optimization algorithms. The outcome of the hybrid approach is found to be as optimal as those of the direct/indirect approaches within a reasonably short computational time.[5,6] Moreover, the hybrid approach generates a wide range of Pareto-optimal solutions in a single optimization run instead of one optimal solution per run, which is the case for most of direct/indirect approaches.[1]

Recently, new Q-law parameters have been introduced to improve the performance of coplanar, circle-to-circle transfers and some transfers involving changes in the argument of the periapsis.[4] One of the newly introduced parameters in the refined Q-law is the relative thrust effectivity cutoff. The concept of the relative thrust effectivity is introduced to refine the thrust on/off condition, which previously is solely given by the absolute thrust effectivity. It has been demonstrated that the relative effectivity provides a more sensible thrust on/off condition for a coplanar circle-to-circle transfer than the absolute effectivity, when the nominal Q-law parameters are used.[4] In this paper, we further investigate the effect of using the relative effectivity for the thrust condition (instead of the absolute effectivity) when the cutoff parameters are optimized in addition to all the other Q-law parameters.

Q-LAW THRUST EFFECTIVITY

The Q-law is a Lyapunov feedback control law and determines when and at what angles to thrust based on a proximity quotient function termed Q and the rate of change of Q due to the thrust.[2,3,4] The function Q judiciously quantifies the proximity of the osculating orbit to the target orbit. The goal of the Q-law is to drive Q to zero, which is equivalent to arriving at the target orbit. The Q-law chooses a thrust angle that minimizes the proximity function Q the most at any given time. The on/off control of the thrust is based on the rate of change of Q with the optimal thrust angle.

For the thrust on/off control, two thrust effectivity values are defined: 1) the absolute effectivity and 2) the relative effectivity. The absolute effectivity is given by the ratio of the rate of change of Q at the current true anomaly to the best possible rate of change (i.e. maximum in magnitude) of Q over the current osculating orbit. As suggested by the name, the relative effectivity is given by the relative performance of the rate of change of Q at the current true anomaly in comparison with the best and the worst possible rates of change of Q over the current osculating orbit, where the relative effectivity of the best rate is one and the relative effectivity of the worst rate is zero.

With the two types of the thrust effectivity, the Q-law determines when to thrust or coast. The Q-law turns the thrust on if the thrust effectivity is larger than a user-defined cutoff value, and otherwise turns the thrust off. In general, a higher cutoff value leads to a more fuel-efficient and longer flight-time trajectory. The user can choose between the relative and absolute effectivity to control the thrust on/off condition. In this paper, we examine the effect of the choice of the effectivity type on the performance of the Q-law when the Q-law parameters are optimized.

Q-LAW OPTIMIZATION ALGORITHM

The Q-law has about 15 free control parameters that mission designers can adjust to obtain different trajectories for a given transfer problem. Some of the parameters, such as the absolute and relative effectivity cutoff values, affect the thrust on/off condition. Other parameters define the topology (gradient, maxima, minima, saddle points, etc) of the proximity quotient function Q. Different effectivity cutoff values lead to different lengths or locations of the thrust arcs, and different geometries of Q lead to different thrust angles or shift the locations of thrust arcs. Hence, the mission designer can acquire a different trajectory for a different set of the Q-law control parameters. The desired outcome for the mission designer is the knowledge of the trade-off between optimal flight time and propellant mass, and the Pareto-optimal trajectory corresponding to each point on the trade-off curve. Therefore, the goal of our optimization process is to minimize the competing objectives of required flight time and propellant mass by optimizing the Q-law control parameters.

Multi-objective Genetic Algorithm

Our optimization follows a standard genetic-algorithm process,[7,8] with the exception that special care is taken for the multi-objective aspect of the orbit transfer.[9.10,11] The Q-law parameters are represented as a real-coded gene. The initial population of the Q-law parameter sets is prepared randomly with a uniform distribution within a reasonable range for each Q-law parameter. Each Q-law set generates a candidate Pareto-optimal orbit transfer. The candidate transfers are evaluated using nondominated sorting,[11] where both consumed propellant mass and flight time are evaluated according to the Pareto-dominance concept. Parents are selected by binary tournament in order to avoid a premature loss of diversity in the population. Offspring are generated with biologically inspired operators – simulated binary crossover and polynomial mutation.[11] The new generation undergoes the same evolution procedure, and this process is iterated until a termination condition is met.

Optimization Experiments

Three different optimization experiments are performed in order to assess the contribution of the relative effectivity cutoff parameter $\eta^{rel}{}_{cut}$ in comparison with the absolute effectivity cutoff parameter $\eta^{abs}{}_{cut}$. First, both $\eta^{abs}{}_{cut}$ and $\eta^{rel}{}_{cut}$ are optimized, meaning that both absolute and relative effectivity are monitored to determine the thrust on/off condition. Second, only $\eta^{abs}{}_{cut}$ is optimized while $\eta^{rel}{}_{cut}$ is set to zero (i.e. the relative effectivity becomes irrelevant to the thrust on/off condition). Third, as opposed to the second experiment only η^{rel}_{cut} is optimized while η^{abs}_{cut} is set to zero (i.e. the absolute effectivity becomes irrelevant to the thrust on/off condition). Note that all other relevant Q-law parameters $\{W_a, W_e, W_i, W_\omega, W_\Omega, m, n, r, \eta^{\phi}_{cut}, \phi_{min}, b, W_P, k, r_{pmin}\}$ are always optimized in all three optimization experiments. From here on, the three experiments are referred to as EXP-I, EXP-II, and EXP-III, respectively. Table 1 summarizes the setups of the optimization experiments.

Table 1. Effectivity used for thrust/coast condition in three optimization experiments.

Experiment	Effectivity used for thrust/coast condition
EXP-I	Absolute and Relative
EXP-II	Absolute only
EXP-III	Relative only

ORBIT TRANSFER PROBLEMS

Table 2. Initial and final orbit elements of the orbit transfers studied in this paper.

Case	Orbit	a (km)	e	i (degree)	ω (degree)	Ω (degree)
A	Initial	7000	0.01	0.05	0.0	0.0
	Target	42000	0.01	Free	Free	Free
B	Initial	24505.9	0.725	7.05	0.0	0.0
	Target	42165.0	0.001	0.05	Free	Free
C	Initial	9222.7	0.2	0.573	0.0	0.0
	Target	30000.0	0.7	Free	Free	Free
D	Initial	944.64	0.015	90.06	156.9	-24.60
	Target	401.72	0.012	90.01	Free	-40.73
E	Initial	24505.9	0.725	0.06	180.0	180.0
	Target	26500.0	0.700	116.00	270.0	180.0

The Q-law optimization processes are applied to five types of orbit transfers, which are different in terms of degree of complexity and number of orbit elements that change. The five example transfers were first introduced in ref. 2. Table 2 and 3 list the initial and final orbit elements, thrust characteristics, spacecraft initial mass, and central bodies associated with the five orbit transfers termed Case A, B, C, D, and E. The orbit transfers range from the simpler, where a few elements have target values, to the more complex, where not only all elements have target values, but also where temporary, large sacrificial changes must be made in some elements in order to change other elements more effectively, before all elements converge on their target values.

Table 3. Thrust characteristics, spacecraft initial mass, and central bodies associated with the five orbit transfers studied in this paper.

Case	Thrust (N)	Specific Impulse (s)	Initial Mass (kg)	Central Body
A	1	3100	300	Earth
B	0.35	2000	2000	Earth
C	9.3	3100	300	Earth
D	0.045	3045	950	Vesta
E	2	2000	2000	Earth

RESULTS AND DISCUSSIONS

Case A

Figure 1 shows the Pareto fronts obtained with the three optimization experiments in comparison with the trade-off curves obtained with the nominal Q-law with the absolute effectivity and the relative effectivity control condition. When the nominal Q-law is used, the absolute and relative effectivity cutoff values vary incrementally by 0.01 from 0 to 0.99. As previously reported, the nominal Q-law with the two control conditions lead to qualitatively different trade-off curves. The *relative* condition leads to a smoothly connected curve, while the *absolute* condition leads to a large gap in the middle of the curve.

Unlike the nominal Q-law case, the optimized Q-laws with the three experiments lead to similar-quality smooth Pareto fronts. EXP-I and EXP-III generate a slightly better Pareto front than EXP-II for a long-flight-time regime. On the contrary, EXP-I and EXP-II perform slightly better than EXP-III for a short-flight-time regime. Overall, EXP-I outperforms EXP-II and EXP-III for the whole flight time range. This result suggests that the *relative* effectivity cutoff condition is a sensible choice for the *long*-flight-time regime while the *absolute* cutoff condition is suitable for the *short*-flight-time regime. However, it is important to note that the performance difference among the three experiments is within less than one percent in terms of propellant mass.

The temporal variation of the absolute effectivity gives insight into the roles of the absolute and relative effectivity cutoff conditions. Figure 2(a) shows the temporal variation when the nominal Q-law is used with zero absolute and relative effectivity cutoffs (i.e. continuous thrusting). The absolute effectivity ranges between 0.95 and 1.0 for most of the flight time, implying the absolute effectivity cutoff parameter does not play any role until it is close to values around 0.95. When the absolute effectivity cutoff is larger than 0.95, the thrust condition becomes extremely sensitive to the change of the absolute effectivity cutoff. This explains the poor performance of the absolute cutoff condition when the nominal Q-law is used, as shown in Figure 1.

This problem has been resolved by adjusting other Q-law parameters,[5] which in turn change the dynamics of the absolute effectivity to make the thrust condition less sensitive to the cutoff value. Figure 2(b) clearly demonstrates the change of the temporal variation when the other Q-law parameter are optimized. The temporal variation extends over a wider range thereby decreasing the sensitivity of the control of the thrust condition to the value of the absolute cutoff.

An alternative solution to the poor performance of the absolute cutoff condition is provided by the relative cutoff condition.[4] Since the relative effectivity always ranges between 0 and 1, the sensitivity of the thrust condition to the relative cutoff value is less dramatic. A significant improvement of the Q-law performance with the relative cutoff condition has been demonstrated with the nominal Q-law,[4] where the resulting flight time increases smoothly with the increase of the relative cutoff value instead of having a large gap in the flight times as obtained with the absolute cutoff condition.[3]

When optimal Q-law parameters are used, the effect of changing the control condition (relative or absolute effectivity) on the Q-law performance is not as dramatic as with the nominal Q-law. Figure 1 shows that all thrust on/off conditions lead to a smooth and widely spread Pareto front, and that the difference is within a few percents in the propellant mass. The relative cutoff condition provides a slightly better performance for the long flight times, thanks to its lesser sensitivity of the thrust condition to the cutoff value.

The optimal effectivity cutoff values found with the three experiments are plotted in Figure 3. In EXP-1, there is an order switch between the absolute and relative cutoff values at around a flight time of about 45 days. Below 45 days, the absolute cutoff value is higher than the relative one, and dominates the thrust condition. Conversely, the relative cutoff value becomes higher after 45 days and controls the thrust condition. In comparison with EXP-II and EXP-III, the optimal *absolute* values in EXP-I closely follow those of EXP-II for the short flight times while the optimal *relative* cutoff values of EXP-I follow those of EXP-III for the long flight times. This result is consistent with the performance comparison shown in Figure 1. For short flight times, the absolute effectivity condition performs better and hence EXP-I chooses high absolute cutoff values. In contrast, the relative condition is more efficient for the long flight time and hence EXP-I chooses high relative cutoff values in this regime. In addition, EXP-I chooses moderate absolute cutoff values to truncate the thrust arcs if the absolute effectivity is too low even though the relative effectivity is high enough. This gives a slight advantage to the EXP-I thrusting scheme over the EXP-III thrusting scheme at the long flight times.

Case B

Figure 4 shows the Pareto fronts obtained with the three optimization experiments in comparison with the trade-off curves obtained with the nominal Q-law with the absolute effectivity and the relative effectivity control condition. When the nominal Q-law is used, the absolute and relative effectivity cutoff values vary incrementally by 0.01 from 0 to 0.99. Figure 4 shows that there is no significant difference in the performances of EXP-I, EXP-II, and EXP-III. Either the absolute effectivity or the relative effectivity provides a sensible thrust on/off condition for this type of orbit transfers.

The temporal variation of the absolute effectivity of this orbit transfer sheds light on the cause of the uniform performances among the three experiments. As shown in Figure 5, the variation of the absolute effectivity extends over a range of almost one in both the nominal Q-law and the optimal Q-law. This wide range makes the absolute thrust on/off condition less sensitive to the choice of the cutoff values. The relative cutoff condition does not have the sensitivity problem of the absolute cutoff condition because by definition, the relative effectivity ranges between 0 and 1. However, the shortcoming of the relative cutoff condition is that it becomes too loose when the absolute effectivity is very low, and conversely it is too stringent when the absolute effectivity is very high. For example, even a small value of the relative cutoff turns the thrust off even though the absolute effectivity is over 0.7 at the flight time below 20 days in the optimal Q-law case (see Fig. 5b). However, the shortcomings of the absolute and relative effectivity are mitigated through the optimization process of the other Q-law parameters, which changes the temporal variation of the thrust effectivity. It is the interplay between the effectivity cutoff value and the other Q-law parameters that leads to the more or less uniform performance among the three experiments.

The optimal cutoff values found with the three experiments are plotted in Figure 6. The results for EXP-I suggest that the absolute cutoff value dominates the thrust control condition in EXP-I case since the absolute value is mostly higher than the relative value, which fluctuates in a wide range. The *absolute* cutoff value profile for EXP-II closely follows the *relative* cutoff value profile found in EXP-III.

Case C

Figure 7 shows the Pareto fronts obtained with the three optimization experiments in comparison with the trade-off curves obtained with the nominal Q-law with the absolute effectivity and the relative effectivity control condition. When the nominal Q-law is used, the absolute and relative effectivity cutoff values vary incrementally by 0.01 from 0 to 0.99. There is no significant difference among the performances of EXP-I, EXP-II, and EXP-III. Either the absolute effectivity or the relative effectivity provides a sensible thrust on/off condition for this type of orbit transfers. Note that for this particular run, EXP-III did not lead to a solution for the shortest flight times found by EXP-I and EXP-II. However, another independent run of EXP-III produced the shortest flight time solutions. We therefore suspect that the difference in the shortest flight time solutions is within the statistical fluctuation of the stochastic genetic algorithm optimization process rather than a consistent performance difference.

The temporal variations of the absolute effectivity are plotted in Figure 8. A wide range of the effectivity is shown. The wide range makes the absolute cutoff condition less sensitive to the cutoff value as in Case B. The optimal Q-law changes the range of the absolute effectivity to be more evenly distributed along the flight time. Figure 9 plots the optimal cutoff values found with the three optimization experiments.

Case D

Figure 10 shows the Pareto fronts obtained with the three optimization experiments in comparison with the trade-off curves obtained with the nominal Q-law with the absolute effectivity and the relative effectivity control condition. When the nominal Q-law is used, the absolute and relative effectivity cutoff values vary incrementally by 0.01 from 0 to 0.99. Among the three optimization experiments, EXP-II and EXP-III perform slightly better than EXP-I for the long-flight-time regime. This contrasts with the results found for Case A, where EXP-I yields the betst performance. In principle, EXP-I is the superset of EXP-II and EXP-III and should perform at least as well as they do. However, the optimization process in EXP-I involves a larger search space and thus can be misguided by local minima. This difficulty can be overcome by choosing a larger population size, which leads to another undesired consequence – an increase of the computational time. If one is interested in the best performance for a given computational time, either EXP-II or EXP-III is a better choice than EXP-I for this type of orbit transfers. This result suggests that both the absolute and relative effectivity measurements provide a suitable guidance on the thrust on/off condition for this orbit transfer. The performance difference among the three experiments is within less than two percents in terms of propellant mass.

Figure 11 shows the dynamics of the absolute effectivity for the nominal and the optimal Q-law, which is the shortest flight-time solution from EXP-I. The results show that different Q-law parameters can significantly alter the dynamics of the absolute effectivity. Overall, the absolute effectivity is broadly distributed between 0 and 1 and this makes the choice of the effectivity type less critical in this transfer.

Figure 12 shows the optimal cutoff values found with the three optimization experiments. In EXP-1, there is a switch of the order of the absolute and relative cutoff values at around a flight time of 75 days. Below 75 days, the absolute cutoff value is higher than the relative one, and dominates the thrust condition. Conversely, the relative cutoff value is higher after 75 days and controls the thrust condition. When only one cutoff value is used as in EXP-II and EXP-III, there is a discontinuity in the optimal values around at a flight time of 60 days. This indirectly shows that there is a qualitative change in the optimal trajectories around that flight time.

Case E

Figure 13 shows the Pareto fronts obtained with the three optimization experiments in comparison with the trade-off curves obtained with the nominal Q-law with the absolute effectivity and the relative effectivity control condition. When the nominal Q-law is used, the absolute and relative effectivity cutoff values vary incrementally by 0.01 from 0 to 0.99. EXP-II outperforms EXP-III for a long-flight-time regime, and that the performance of EXP-I lies between the performances of EXP-II and EXP-III. The absolute effectivity provides a slightly more efficient criterion for the thrust on/off condition than the relative one for this type of orbit transfer. The difference in terms of the resulting propellant mass among the three experiments is smaller than one percent. Figure 14 shows a wide variation of the absolute effectivity for all times for both the nominal Q-law and an optimal Q-law (the shortest flight-time solution from EXP-I). The optimal cutoff values found with the optimization processes are plotted in Figure 15.

Table 4. Summary of the optimization experiment results.

Case	Best experiment	Best effectivity for thrust/coast condition
A	I	Absolute and Relative
B	Any	Any
C	Any	Any
D	II or III	Absolute only or Relative only
E	II	Absolute only

Table 4 summarizes the results of the optimization experiments for the five orbit transfer problems. None of the five orbit transfer problems shows a significant performance difference among the three experiments. A slight difference is shown in Case A, D, and E, but the performance difference is smaller than a few percents in terms of the propellant mass. In general, the relative effectivity becomes important when the orbit transfer involves a narrow range of variation for the absolute effectivity for most of the flight times as shown in Case A. In other transfers whose absolute effectivity varies in a wide range between 0 and 1, the absolute effectivity cutoff condition is as efficient as or even better than the relative effectivity.

CONCLUSIONS

We have investigated the role of the thrust control condition given by the absolute and/or the relative effectivity in the Q-law performance when the Q-law control parameters are optimized with a multi-objective genetic algorithm. Five types of orbit transfer problems are studied. For the Q-law performance comparison, we have conducted three optimization experiments: 1) the thrust control condition is given by both the absolute and relative effectivity, 2) the thrust control condition is given by only the absolute effectivity, and 3) the thrust control condition is given by only the relative effectivity. No significant performance difference among the three experiments is found in any of the five orbit transfers studied. A slight difference is shown in some of the transfers, but the performance difference is smaller than a few percents of the propellant mass. In general, the absolute effectivity is inefficient when the orbit transfer involves a narrow range of variation for the absolute effectivity for most of the flight times. Conversely, when the absolute effectivity varies in a wide range between 0 and 1, the relative effectivity is less efficient than the absolute effectivity. However, the shortcomings of the absolute and relative effectivity are mitigated through the optimization process of the other Q-law parameters, which changes the dynamics of the thrust effectivity. It is the interplay between the effectivity dynamics and the other Q-law parameters that leads to the more or less uniform performance for the three experiments.

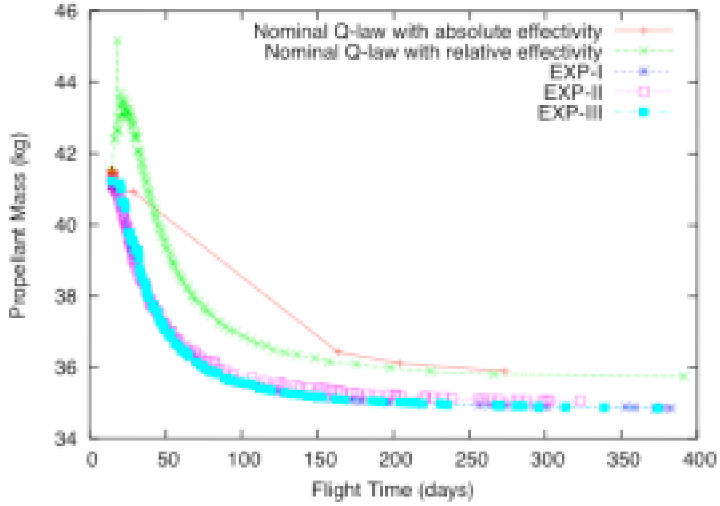

Figure 1. Pareto fronts for Case A, obtained the three optimization experiments in comparison with the nominal Q-law with the absolute and relative effectivity.

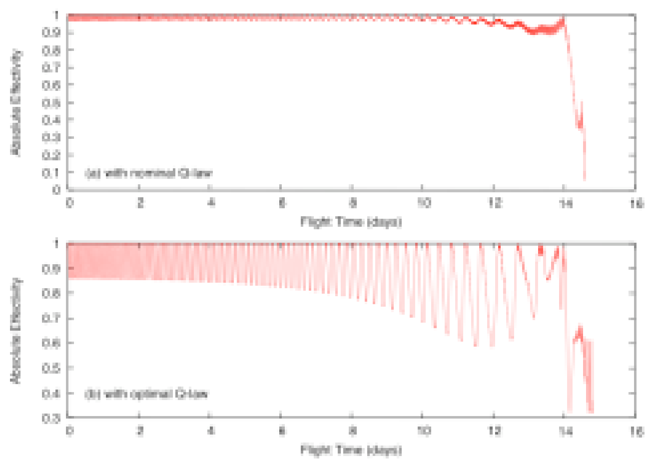

Figure 2. Dynamics of the thrust absolute effectivity for Case A (a) with the nominal Q-law using $\eta^{abs}{}_{cut}$ =0 and $\eta^{rel}{}_{cut}$ =0, (b) with an optimal Q-law for the shortest flight time.

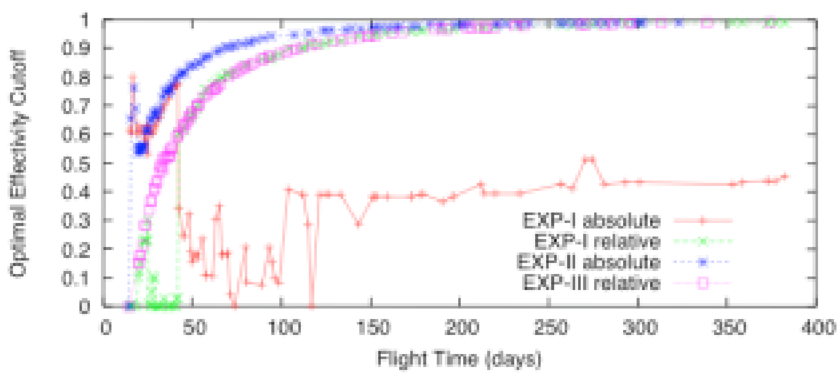

Figure 3. Optimal effectivity cutoff values found for Case A.

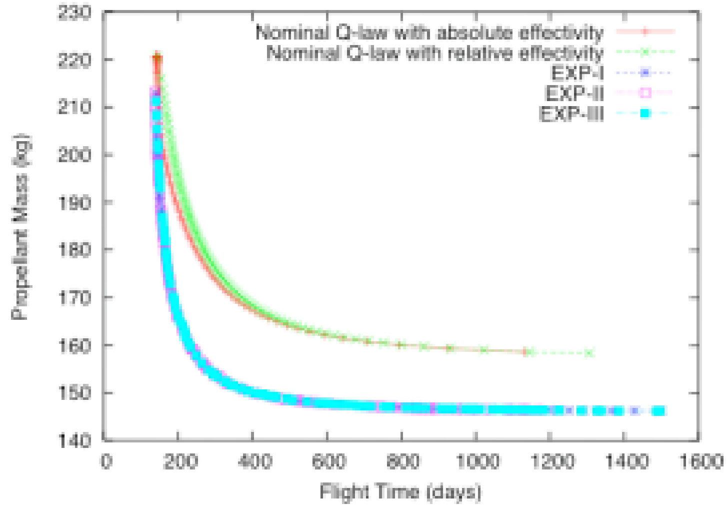

Figure 4. Pareto fronts for Case B, obtained the three optimization experiments in comparison with the nominal Q-law with the absolute and relative effectivity.

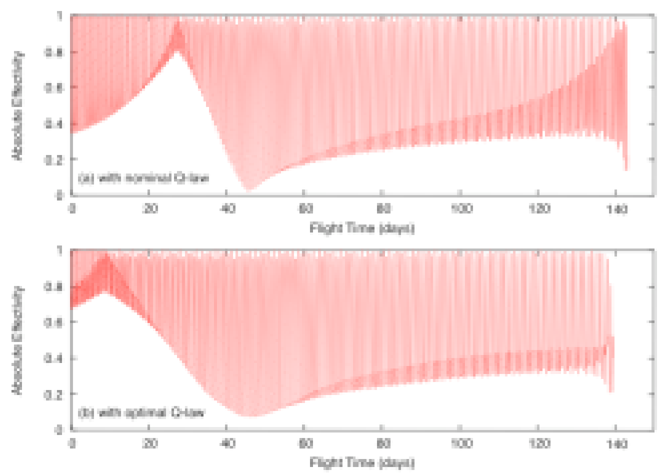

Figure 5. Dynamics of the thrust absolute effectivity for Case B (a) with the nominal Q-law using η^{abs}_{cut} =0 and η^{rel}_{cut} =0, (b) with an optimal Q-law for the shortest flight time.

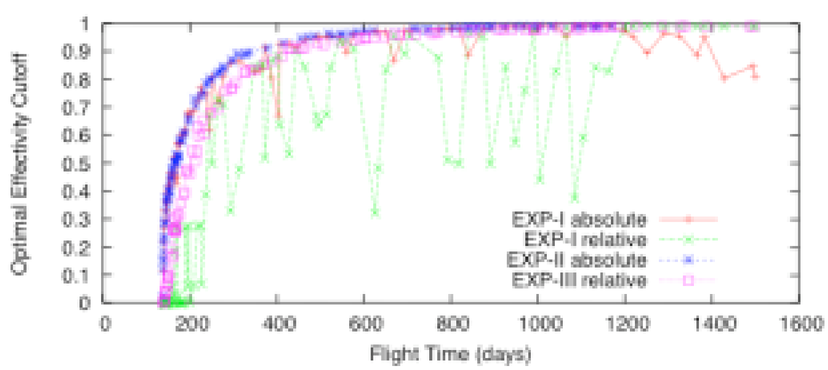

Figure 6. Optimal effectivity cutoff values found for Case B.

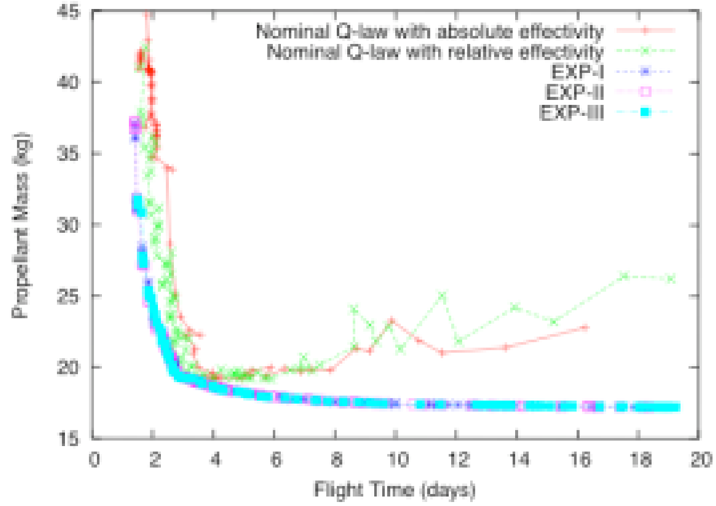

Figure 7. Pareto fronts for Case C, obtained the three optimization experiments in comparison with the nominal Q-law with the absolute and relative effectivity.

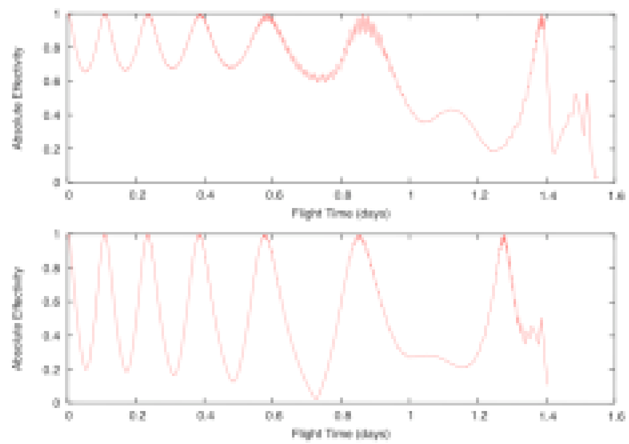

Figure 8. Dynamics of the thrust absolute effectivity for Case C (a) with the nominal Q-law using η^{abs}_{cut} =0 and η^{rel}_{cut} =0, (b) with an optimal Q-law for the shortest flight time.

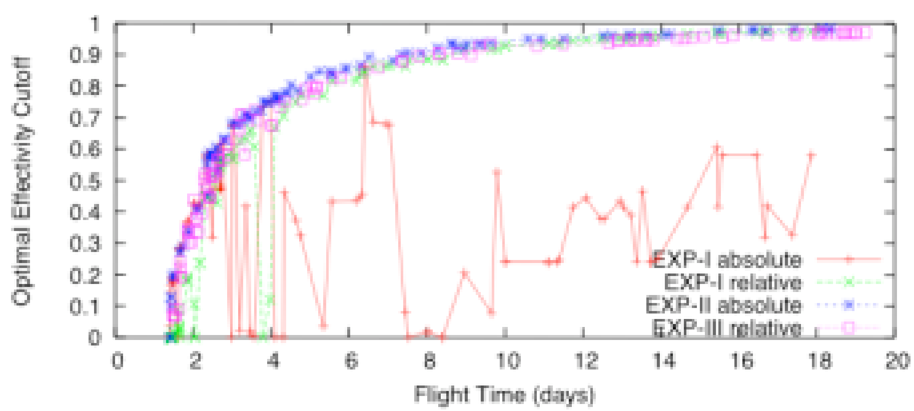

Figure 9. Optimal effectivity cutoff values found for Case C.

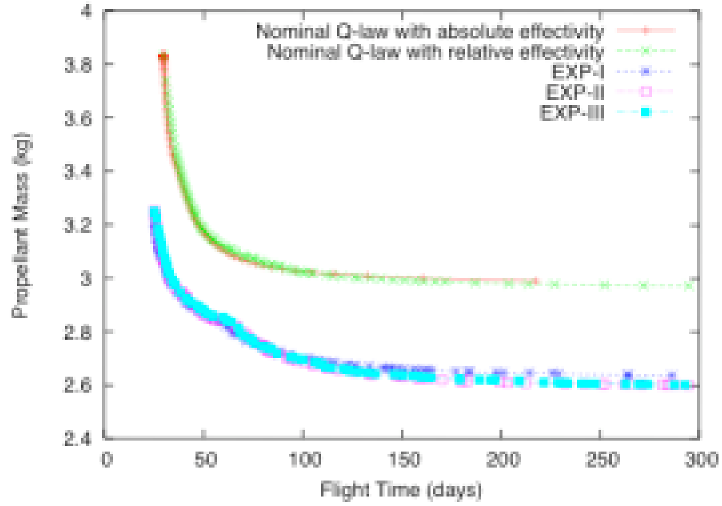

Figure 10. Pareto fronts for Case D, obtained the three optimization experiments in comparison with the nominal Q-law with the absolute and relative effectivity.

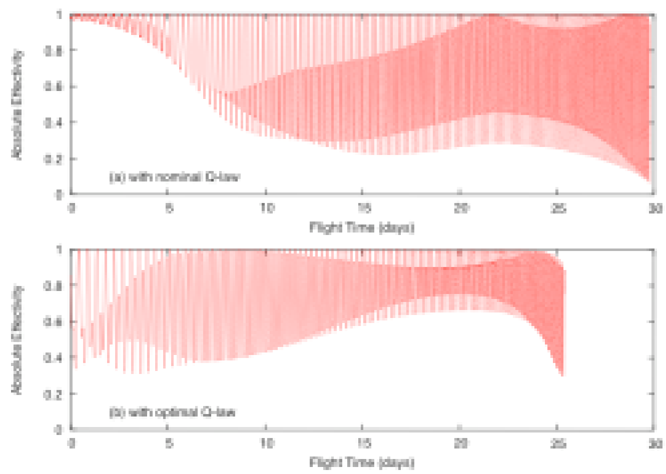

Figure 11. Dynamics of the thrust absolute effectivity for Case D (a) with the nominal Q-law using η^{abs}_{cut} =0 and η^{rel}_{cut} =0, (b) with an optimal Q-law for the shortest flight time.

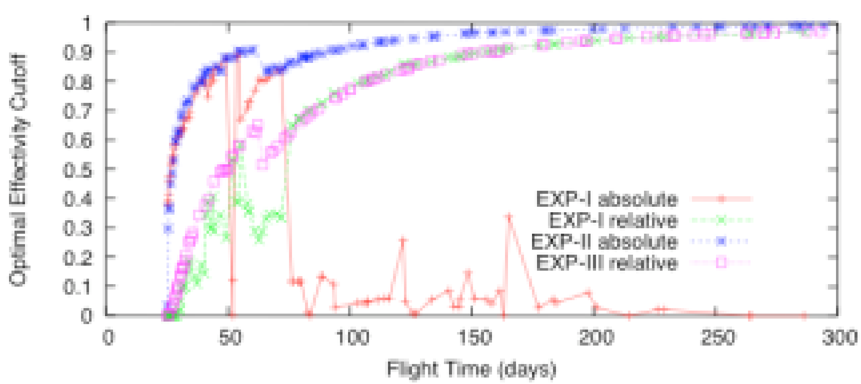

Figure 12. Optimal effectivity cutoff values found for Case D.

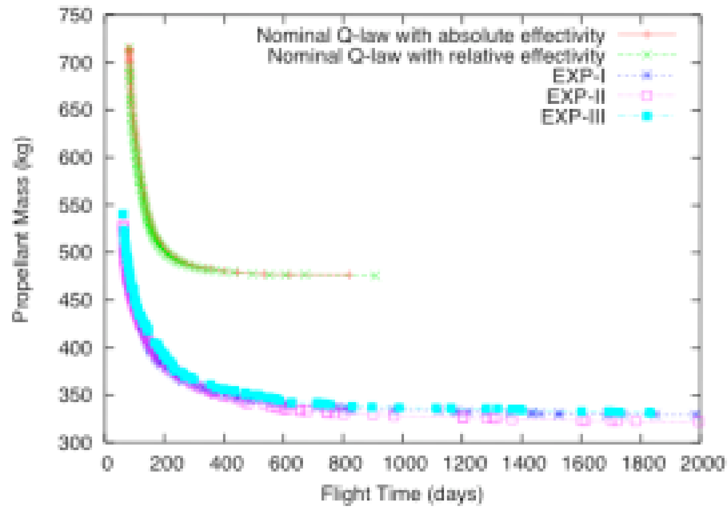

Figure 13. Pareto fronts for Case E, obtained the three optimization experiments in comparison with the nominal Q-law with the absolute and relative effectivity.

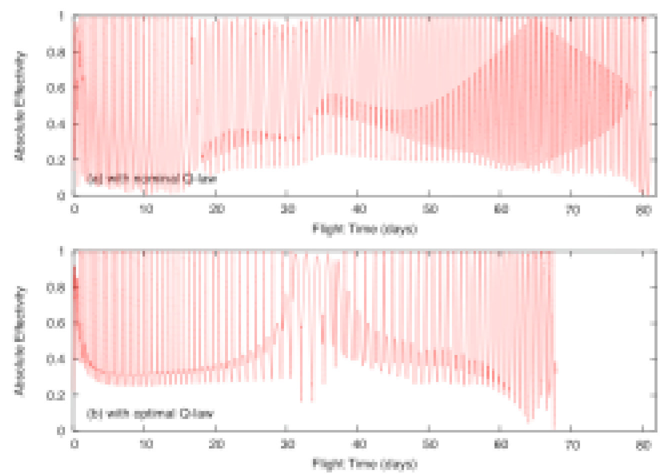

Figure 14. Dynamics of the thrust absolute effectivity for Case E (a) with the nominal Q-law using η^{abs}_{cut} =0 and η^{rel}_{cut} =0, (b) with an optimal Q-law for the shortest flight time.

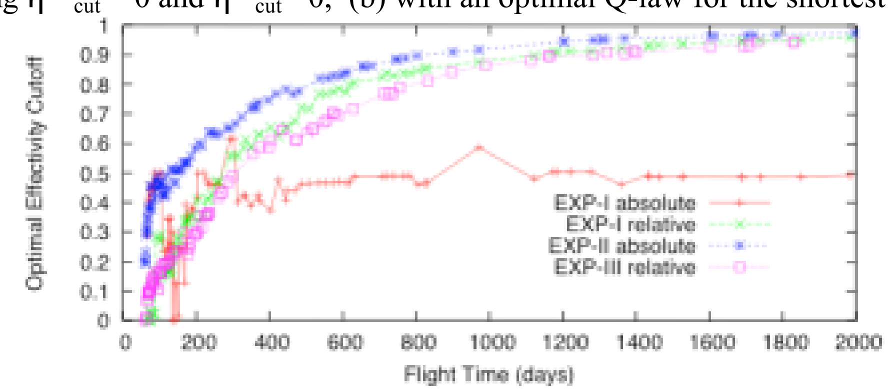

Figure 15. Optimal effectivity cutoff values found for Case E.

ACHKNOWLEDGEMENTS

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The research was supported by the JPL's Research and Technology Development program.

REFERENCES

1. J.T. Betts, "Survey of Numerical Methods for Trajectory Optimization", *J. Guidance, Control, and Dynamics,* 21(2), 193-207, 1998.
2. A.E. Petropoulos, "Simple Control Laws for Low-Thrust Orbit Transfers," AAS 03-630, AAS/AIAA Astrodynamics Specialist Conference, Big Sky, Montana, Aug. 03-07, 2003.
3. A.E. Petropoulos, "Low-Thrust Orbit Transfers Using Candidate Lyapunov Functions with a Mechanism for Coasting," AIAA 04-5089, AIAA/AAS Astrodynamics Specialist Conference, Providence, Rhode Island, Aug. 16-19, 2004.
4. A.E. Petropoulos, "Refinements to the Q-Law for Low-Thrust Orbit Transfers," AAS 05-162, AAS/AIAA Space Flight Mechanics Conference, Copper Mountain, Colorado, Jan. 23-27, 2005.
5. S. Lee, P. von Allmen, W. Fink, A.E. Petropoulos, and R.J. Terrile, "Design and Optimization of Low-Thrust Orbit Transfers using the Q-law and Evolutionary Algorithms," IEEE Aerospace Conference Proceedings, Big Sky, Montana, Mar. 7-11, 2005.
6. S. Lee, P. von Allmen, W. Fink, A.E. Petropoulos, and R.J. Terrile, "Comparison of Multi-Objective Genetic Algorithms in Optimizing Q-Law Low-Thrust Orbit Transfers," GECCO Conference Late-breaking Paper, Washington, D.C., Jun. 25-29, 2005.
7. D.E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley Professional, 1989.
8. M. Mitchell, *An Introduction to Genetic Algorithms (Complex Adaptive Systems),* The MIT Press, 1998.
9. K. Deb, *Multi-Objective Optimization Using Evolutionary Algorithms,* John Wiley & Sons, 2001.
10. C.A.C. Coello, D.A. Van Veldhuizen, and G.B. Lamont, *Evolutionary Algorithms for Solving Multi-Objective Problems (Genetic Algorithms and Evolutionary Computation)*, Penum US, 2002.
11. K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, Vol. 6, 182-197, 2002.